Published in final edited form as:

IEEE Spectr. 2019 February ; 56(2): 36–41. doi:10.1109/mspec.2019.8635815.

The Stethoscope Gets Smart:

Engineers from Johns Hopkins are giving the humble stethoscope an AI upgrade

Mounya Elhilali, James E. West

School of Computer Science and Technology, Heilongjiang University, Harbin, China

You wake up one morning to discover that your child is ill: His forehead feels hot to the touch, and his rapid breathing has a wheezing sound. You live in Malawi, where your health care options are few. When the local clinic opens, you wait for your turn with the solitary clinic worker. She's not a doctor, but she's been trained to identify and handle routine problems.

She puts on a stethoscope and presses its chest piece against your son's front and back to carefully listen to his lungs. Through the windows, open in the heat of the day, come the sounds of people talking, the thrum of a generator, and the roar of a moped on the main road. The health worker strains to hear. Will she detect the signs of bacterial pneumonia, give you antibiotics, and recommend that you travel to a hospital in the nearest city? Or will she dismiss the symptoms as signs of a common cold, and send you back home with your son? If she mistakenly sends you home, will your son survive the mistake?

Worldwide, more kids die of pneumonia and other lung ailments than from any other cause. These acute lower respiratory infections kill nearly 1 million children each year worldwide, causing more deaths than HIV and malaria combined. But fewer than 5 percent of people in the developing world have access to the X-ray imaging that's considered the gold standard for pneumonia diagnoses. What's more, many sick children never get to a clinic because they live far away from a facility, roads are in poor condition, and transportation is expensive. In sub-Saharan Africa, where most pneumonia deaths occur, only two in five children with pneumonia symptoms receive professional medical care.

Given the global burden of these respiratory ailments, the World Health Organization has developed strategies for countries with limited resources. The guidelines for pneumonia diagnosis minimize dependence on any technological tools, and instead rely solely on observed symptoms of shortness of breath, cough, and rapid breathing. In hopes of saving lives, the WHO recommends antibiotic treatment for all children with these symptoms, with the result that half the children who get treated for pneumonia don't really have it. This approach puts unnecessary costs on communities, and the unnecessary medication contributes to the growing problem of antibiotic-resistant bacteria.

We propose a technological solution that builds on the stethoscope, which has hardly changed since its invention in the early 1800s. This ubiquitous clinical tool has a number of limitations. For best results, the user should be in a quiet environment, as background noise can easily mask the subtle sounds coming from the lungs. The diagnostician must be well trained in positioning the chest piece properly on the body and in interpreting the sounds,

which is something of an art form. Learning to do it well requires time and mentorship, and even experts can disagree on what they hear. That's why even leading doctors prefer to complement their use of the stethoscope with other techniques, such as chest X-rays.

We wondered whether we could overcome the stethoscope's limitations and thereby take advantage of its low cost and ease of use. With funding from the National Heart, Lung, and Blood Institute, NASA, and the Bill & Melinda Gates Foundation, we set out to completely reengineer the stethoscope.

The project was a collaboration among engineers, doctors, and public health experts at Johns Hopkins University, in Baltimore. Together, we invented a device that uses digital sensing technology for sound capture, active acoustics for noise cancellation, and artificial intelligence (AI) to help health workers make accurate pneumonia diagnoses. We hope that our smart stethoscope will be deployed around the world to prevent children from dying of pneumonia. And as we test its capabilities, we may find that it's a boon not only for the developing world but also for hospitals in the United States and other well-off countries.

Our mission sprang from an ambitious global study of pneumonia that began in 2008, with funding from the Gates Foundation. The Pneumonia Etiology Research for Child Health (PERCH) study explored the causes of pediatric pneumonia across multiple countries in Africa, Asia, and South America.

In the course of this study, a team of doctors from Johns Hopkins set out to determine just how useful lung sounds could be in diagnosing pneumonia. At first, the researchers equipped local health workers with commercially available stethoscopes that they had augmented with small digital recorders. The health workers recorded lung sounds from the children they encountered in rural clinics and urban hospitals, then transmitted the files to expert pulmonologists at Johns Hopkins, who listened to these sounds and offered more informed medical opinions.

Quickly, the use of the commercial stethoscope became problematic. To get a decent recording from the lungs, the health worker had to position the device's chest piece in just the right places on the child's chest and back. Unfortunately, the local workers typically had only basic training in how to use the tool, and often missed the "sweet spots" that best transfer the lung sounds to the stethoscope head. Furthermore, they were usually making their recordings in noisy settings, nothing like the quiet doctor's office for which the stethoscope is intended. Background sounds of people, machines, cars, and motorcycles would mix with the lung sounds, leaving the expert with a recording of cacophony. In addition, the local workers needed Internet access to transmit their files, and connectivity limitations and glitches meant that it often took many hours or even days to get back an expert diagnosis.

If they wanted to get useful recordings, the Johns Hopkins doctors decided, they'd have to entirely rethink the stethoscope. So they approached the university's department of electrical and computer engineering, where one of us (West) is a professor, and the other (Elhilali) is an associate professor. As both of us specialize in audio engineering, we were eager to take up the challenge.

Here's how a classic stethoscope works. Sound is produced by an organ in the body, and these acoustic waves cause a vibration in the stethoscope's chest piece, which acts as a resonator. This chest piece has two sides with carefully chosen shapes—a flat, disk-like diaphragm and a hollow cup called a bell—that naturally oscillate with different frequency ranges, and which are used for different diagnostic tasks. The acoustic vibration travels via an air-filled tube that connects to two ear-pieces, thus relaying the sounds of the patient's body to the listener. Although the stethoscope is designed to maximize sound pickup and delivery to the user's ears, sound levels are usually quite low.

More-modern electronic stethoscopes convert the acoustic wave into electric signals that can be processed in the device to amplify the sounds. They replace the double-sided chest piece with one gadget that uses transducers to convert acoustic signals to electric ones with minimal distortion or noise effects. Users can toggle between “diaphragm” and “bell” modes, and can often adjust the sounds that reach their ears in other ways.

The Johns Hopkins device is an electronic stethoscope that improves on digital devices currently on the market. The upgrades start with its hardware: The chest piece is packed with transducer arrays to achieve a uniform sensitivity over the entire active area. This design delivers a strong signal even when the chest piece isn't placed in precisely the right location, which gives a major assist to untrained users. We've experimented with a variety of transducers, including microelectromechanical systems (MEMS), which allow us to pack a great many microphones into a small area, as well as nanofiber materials that match the acoustic properties of skin to limit signal loss.

For a stethoscope to be useful in clinics in the developing world, dealing with noise is of utmost importance. In classic stethoscopes, environmental noise contaminates the signal at three points: at the chest piece, through the rubber hose, and at the ear of the user. The Johns Hopkins design mitigates noise by improving the coupling between the patient's body and the chest piece, swapping the rubber hose for an electric cable, and employing digital noise-control techniques to ensure that a strong signal arrives at the user's ears. Our stethoscope also has an external microphone in the chest piece to collect ambient sounds. This data can then be used to reduce noise through adaptive signal analysis—as we'll explain in the next section.

Our stethoscope gets its smarts from an onboard microprocessor that serves as an operating system, allowing us to develop various applications to tailor the device for different needs. So far, our team has focused on two apps that we deemed crucial for developing world clinics. One dials down the level of noise; the other dials up the level of expertise.

Reducing external noise is a complicated matter. It's relatively easy to design noise-canceling software that recognizes the signal pattern created by normal breathing and filters out everything else. But pathological conditions like pneumonia create abnormal signal patterns—and an overactive noise-canceling filter can remove that critical information. For example, if a child is wheezing, producing high-pitched whistling sounds in addition to the normal sounds of breathing, the “de-noising” algorithm should not mistake this whistling for high-pitched ambient noise, such as a crying child would make. Conversely, trying

to preserve abnormal patterns can also preserve ambient noise. Our app aims to suppress undesired noise while maintaining the integrity of the signal coming from the patient's lungs.

As we designed our noise-cancellation algorithm, we considered two challenges. First, the lung sounds that indicate possible pathologies are unpredictable and irregular. Second, the noises in a busy clinic are highly variable: The exam room could be filled with chattering people, ringing phones, and whirring fans; the child being examined might be crying out and thrashing around on the exam table; and the movements of the stethoscope itself can produce electronic noise.

With those challenges in mind, our team created a noise-cancellation algorithm that uses active and adaptive acoustics. Instead of just setting up a filter and letting it run in a passive way, our system analyzes both lung sounds and ambient sounds recorded by the external microphone, looking at their frequencies on the audio spectrum. Then the app adapts the degree and spectral span of noise cancellation as the algorithm tracks the two signals over time. We validated this approach through a blind listening test, in which we asked a panel of pediatric pulmonologists to compare the quality of sounds coming from a commercial stethoscope used in a rural clinic with the same signals run through the Johns Hopkins algorithm. The experts preferred our de-noised signal 95 percent of the time.

We've also extended the audio software to deal with extraneous sounds from the body. When a health worker is trying to listen to a patient's lungs, the sound of a beating heart can be distracting and can mask the lung sounds. Pulmonologists train their ears to ignore the heartbeat and focus on the breathing patterns—and we gave our algorithm similar training. Our app cancels out the sound of the heartbeat to provide the user with the pure and unadulterated signal from the lungs.

To save the lives of children throughout the developing world, we wish we could send expert physicians to staff every remote clinic and overcrowded hospital. So we did the next-best thing and developed an app to send physician expertise to these settings.

This second app, the one that makes the Johns Hopkins stethoscope truly smart, automatically screens for cases of pneumonia by distinguishing normal from abnormal breathing patterns, looking in particular for the wheezing or crackling sounds that can indicate liquid and inflammation in the lungs.

Our AI approach was inspired by the brain's auditory system, which analyzes both the overall audio signal and its specific features. The brain devotes the most processing power to the frequencies that change over time and those that it judges most relevant to the situation. Think about being at a noisy cocktail party, and how you're able to ignore the background buzz of chatter to focus on the words spoken by the person right in front of you. As our app analyzes the de-noised lung sound, it takes in the whole signal, but zeroes in on the specific modulations that indicate a case of pneumonia.

We used a machine-learning approach to train the app, giving it plenty of data and a mandate to pick out the most useful features. The algorithm trained on a database of recordings

of about 1,500 patients from five African countries and two Asian countries. When we tested the AI-enabled app, we found that it could automatically distinguish healthy people from those with pneumonia with an accuracy of 87 percent, far surpassing other automatic diagnosis methods. Our team is continuing to work on this algorithm, aiming to improve its performance with machine-learning techniques.

Health workers in rural clinics could use our smart stethoscope even if they don't have access to the Internet; the onboard technology does all the processing and provides an instant recommendation on diagnosis using a small built-in LED display. We envision that the final version of the smart scope will connect via Bluetooth to the user's phone or tablet, where the user can get more information about the diagnosis, visualize the lung-sound signal, and play back previously recorded sounds.

As our smart stethoscope has a programmable platform, we're now working to develop additional apps that would make it a multipurpose tool, useful for many diseases and scenarios. We think it could be useful in diagnosing heart failure and various gut disorders, and we're collaborating with cardiologists and gastroenterologists to collect training data and design algorithms.

Today's electronic stethoscopes typically cost around US \$500, making them far too expensive for many health workers in the developing world. The Johns Hopkins smart scope is designed to be significantly cheaper, with affordable electronic components and low-cost power and computing options, in hopes that it can be useful for low-resource communities.

Right now, our smart scope is undergoing extensive field testing. Clinic workers in Peru, Bangladesh, and Malawi are trying out prototypes, as are doctors at the Johns Hopkins pediatric emergency room. These testing and validation studies are assessing how well it functions in noisy settings and how accurate its pneumonia diagnoses are across a diverse population of patients. In these studies, we're evaluating the lung sounds using objective signal quality metrics, and we're validating the diagnoses by comparing them with diagnoses made via chest X-rays and the listening assessments of pulmonology experts.

While we continue our work in academia, some of our colleagues have founded a startup called Sonavi Labs to quickly turn our prototype into a commercially viable product. The mission of Sonavi Labs is to develop modern medical devices that effectively use lung sounds for the diagnosis of disease. Their first products, digital stethoscopes called the Felix and FelixPro, are set to launch in early 2019.

By enabling on-the-spot screenings with a cheap device that can be handled by local health workers, we hope our reengineered scope will make a difference in the global health crisis of childhood pneumonia. Viewed on the macro level, we hope to scale up our efforts so that hundreds of thousands of lives can be saved. But on the micro level, we feel that if even one parent sees her child recover because of early diagnosis, all our hard work will have been worth it.

AI ASSISTANCE:

A health worker in Bangladesh listens to sounds from a boy's lungs with the help of the Johns Hopkins smart stethoscope, which cancels out extraneous noise and identifies telltale signs of pneumonia.

GLOBAL KILLER:

In 2015, about 5.9 million children died before the age of 5. Pneumonia accounted for 15.5 percent of those deaths around the world, and was the top cause of death in sub-Saharan Africa.

SOURCE: THE LANCET, VOLUME 388, ISSUE 10063, DECEMBER 2016

SOUNDS OF TROUBLE:

These two spectrograms show digitized audio signals from a child who's breathing normally [left] and from a child who's producing the high-pitched wheezing noises [right] that indicate fluid and inflammation in the lungs.

WORKING PROTOTYPES:

After an early version of the smart stethoscope [top] proved effective at filtering out noise and giving the listener a clear signal from the lungs, Johns Hopkins researchers developed the current version [bottom], which includes an LED display that gives the user a recommendation on the diagnosis.

Ambient noise

[1] The smart scope has an external microphone that records ambient sounds. The system uses active and adaptive acoustics to remove those sounds, in real time, from the digitized audio signal transmitted to the user's ears.

Noisy body sounds

[2] Pulmonologists are trained to ignore other body sounds, such as the heartbeat, when they listen to the lungs. The smart scope detects those extraneous sounds and removes them, too.

Enhanced signal

[3] What's left is a clear audio signal containing only the sounds from the patient's lungs. With this refined signal, both the listener and the stethoscope's software can identify the signs of pneumonia.

SEEKING CLARITY:

Traditional stethoscopes are designed for quiet doctors' offices, but many health clinics around the world are far from quiet. The devices are also meant to be used by experts trained in the subtle art of interpreting body sounds. The smart stethoscope provides a clear audio signal by canceling out both noise from the environment and irrelevant noise from inside the body.